LAWRENCE LIVERMORE NATIONAL LABORATORY

LLNL-PROC-408339

Defect Classification Using Machine Learning

A. Carr, L. Kegelmeyer, Z. M. Liao, G. Abdulla, D. Cross, W. P. Kegelmeyer, F. Raviza, C. W. Carr

October 31, 2008

Boulder Damage Symposium
Boulder, CO, United States
September 22, 2008 through September 24, 2008

Defect Classification Using Machine Learning

Adra Carr*, L. Kegelmeyer, Z. M. Liao, G. Abdulla, D. Cross, W.P. Kegelmeyer, F. Ravizza, and C. W. Carr**

Lawrence Livermore National Laboratory, 7000 East Ave., L-592, Livermore, CA 94550

* University of Colorado at Boulder, Boulder, CO 80305

** Sandia National Laboratory, Livermore, CA 94550

ABSTRACT

Laser-induced damage growth on the surface of fused silica optics has been extensively studied and has been found to depend on a number of factors including fluence and the surface on which the damage site resides. It has been demonstrated that damage sites as small as a few tens of microns can be detected and tracked on optics installed a fusion-class laser, however, determining the surface of an optic on which a damage site resides in situ can be a significant challenge. In this work demonstrate that a machine-learning algorithm can successfully predict the surface location of the damage site using an expanded set of characteristics for each damage site, some of which are not historically associated with growth rate.

Keywords: laser-induced damage, machine learning, laser damage surface dependence, laser damage growth, growth modeling

1. INTRODUCTION

Precise models of laser damage growth are necessary to accurately predict how and when a damage site will require its host optic to undergo maintenance. Laser-induced damage growth on the surface of fused silica optics has been extensively studied and has been found to depend on a number of factors including fluence, pulse duration, wavelength and the surface which the damage site resides. Of these factors, the total fluence incident on the site and the surface of the site (input or exit) are two of the most important and are needed to make the most accurate predictions of future growth. Sites on the exit surface are observed to grow exponentially with laser fluence while sites on the input surface are observed to grow linearly [1, 2]. As a result, it is critical to be able to determine a damage site's surface for damage site growth prediction.

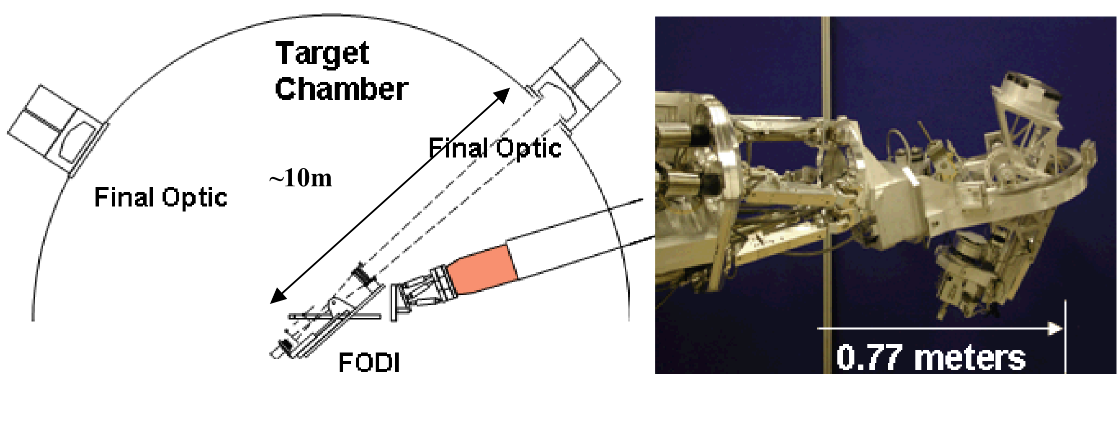

Figure 1: Schematic of where the FODI image system positions within the target chamber (left), with a distance ~10m from the final optics assembly and a close up of the FODI image system itself (right).

At the National Ignition Facility (NIF), a sophisticated online inspection system, FODI (Final Optic Damage Inspection) is used to detect the initiation of damage sites and track their growth [3]. FODI records light scattered off of defects on each optic under either edge or back-illumination. The system is able to detect damage sites greater than ~50μm in diameter with a confidence of ~99% and measurement error (1σ) ≤ ± 15%. For a given detected damage site, post processing using custom image analysis tools [4] can determine a number of aspects of a site's recorded image (i.e. position, size, eccentricity, Signal to Noise (S/N) of the scattered signal, etc). However, due to the large working distance (and hence depth of focus) of the FODI camera and the relative thinness of each optic, FODI has limited effectiveness in determining a site's surface by best focus, as illustrated with typical image capture by FODI of an input and exit surface site (Figure 2).

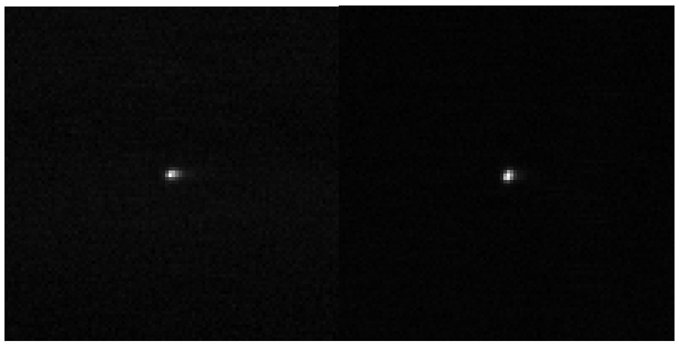

Figure 2: Example of site capture by online image system. Left image is an exit surface site of ~40 μm and right image is an input site of ~50 μm.

2. ANALYTICAL RESULT

One way of determining the surface on which a defect resides is to exploit the difference in growth rates of input vs. exit surface sites. The mean growth for an exit surface site can be written as:

$$D = D_0 e^{0.038\cdot(\phi-5)\cdot\Phi(\phi-5)} \qquad (1)$$

where D is the predicted size of the site, D_0 is the current site size, and ϕ is the total fluence the site is exposed to and Φ being the Heavyside function. The growth of input surface site can be written as:

$$D = D_0\left(0.009\cdot(\phi-4)\cdot\Phi(\phi-4)\right) \qquad (2)$$

Analytically, one can construct an algorithm to classify on which surface a damage site is situated by:

1. Calculating the expected growth with the input and exit surface models
2. Comparing with measured growth coefficient and assign to the surface with the closest value to measured growth

Measuring the *actual* surface a site resides on using a microscope (after the optics have been removed from they system), one can compare the accuracy of the analytical classification scheme to the truth data. However, the growth behavior of sites on each surface overlaps sufficiently (see Figure 3), making this classification scheme ineffective. Due to this overlap, predicting site surface based on growth coefficients yielded 35.5% accuracy. Alternatively, sites can be classified biased on size, with the expectation that exit surface sites will, in general, grow faster and therefore be larger (see Figure 3). Due to the low accuracy in determining site surface based on growth coefficient or size, neither attribute was seen to be a good metric in determining site surface by itself.

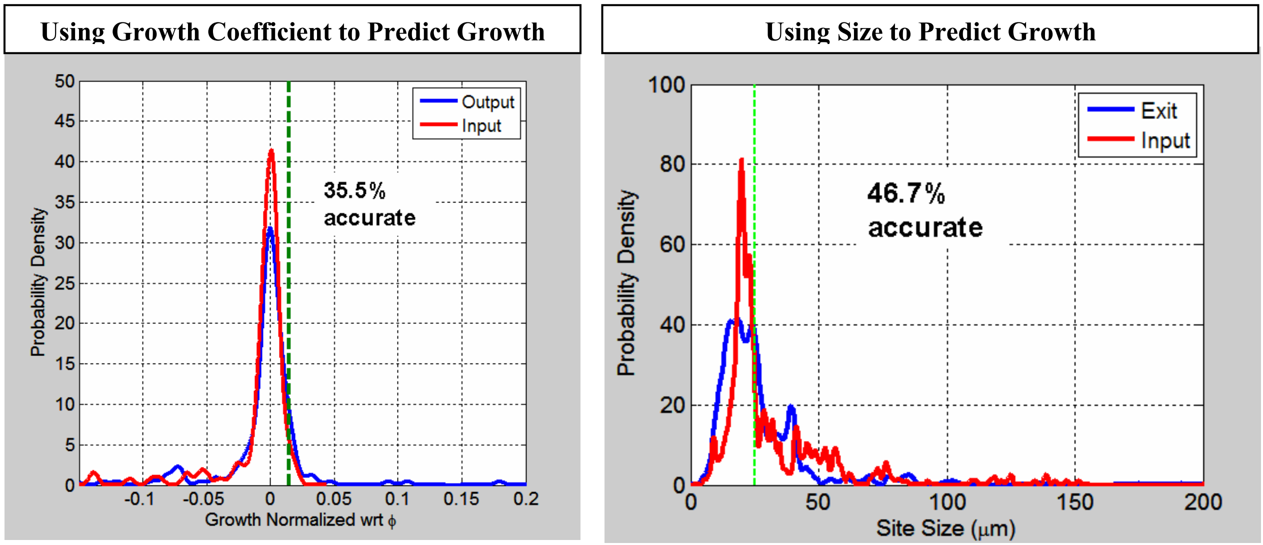

Figure 3: Normalized growth coefficients (to fluence) of damage sites on the input vs. exit surface (left). Size distribution of input vs. exit surface (right).

3. MACHINE LEARNING

A suite of machine learning software, Avatar Tools [4] was used to infer a site's surface based on the previously mentioned data. Currently, Avatar is used on NIF to determine if a detected defect is an actual damage site, a phase object, or a camera flaw, etc. It is able to make these classifications with ~99% accuracy [5]. One of the main advantages of machine learning is that a computer is not constrained by *n*-space parameters as humans typically are. One representation of how Avatar makes its inferences/conclusions about a site's surface is by creating a feature space composed of all the attributes associated with a given data point (a given damage site). Figure 4 shows a plot of a fictional 3-dimensional feature space with example attributes. One can then partition this feature space based on the site's known surface. Therefore, plotting new points (new damage sites) in this feature space with no known surface information, each point will fall into a certain segmented space, thus allowing an inferred conclusion of whether a site is on the input/exit surface. Overall, Avatar "learns" these classifications (segmentations of feature space) by:

1. Using an expert-labeled (In this case known input/output surface from off-line microscope measurements) dataset as "training data"
2. Using all attributes of the training data to group/classify data into different attribute partitions spaces – specifically, by building a forest of decision trees using combinations of the attributes/parameters
3. The majority vote from this forest of decision trees which arrive at the correct answer as defined by the training data is used to classify new data points that lack surface information and infer or "guess" for which surface the site is located.

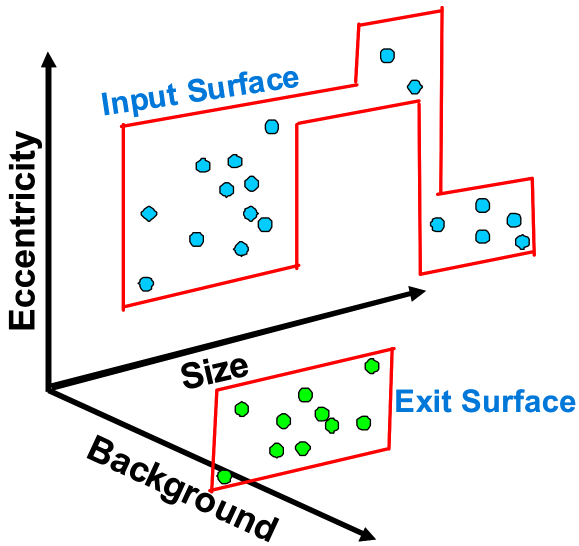

Figure 4: Example of fictional feature space (3D) based on the size, the eccentricity of the site, and the background value of the site. In this example, we can segment the feature space based on whether the site is on the input or exit surface of an optic.

4. RESULTS

Avatar was "trained" using roughly 1500 data points (~ 60 unique sites each with over 20 unique shots on each site) on a fused silica optic with 27 different attributes/parameters associated with each site. A selected number of these attributes are listed in Table 1. Note that no fluence information of any kind was provided to Avatar. This ensured that the data used to train Avatar only depended on the physical properties of a damage site, all of which would be available for a site while the optics is still in situ on a NIF beamline. Using the classifications that Avatar "learned" from a subset of data (for training), we could cross-validate Avatar's inferred conclusions. This produces an overall *confusion matrix* as illustrated in Figure 5. Since we are working with sites with known surface information, this matrix tells how many damage sites' surface designations were misclassified from a portion of the data. The Avatar was able to produce an overall prediction accuracy of 95.3%.

Table 1: **Selected attributes from the 27 attributes associated with each damage site. Most attributes are determined from the original FODI scattering image.**

Description Name	Description of Attribute
AREA_IN_PIXELS	Area in pixels of the measured site
SNR	signal to noise ratio based on the peak pixel value and the standard deviation of the background
RGP	Mean of the intensity values of the perimeter within the gradient image
RGP_SD	Standard deviation of the intensity values of the perimeter within the gradient image
RGP_NORM	Normalization factor used with RGP
LAXIS	The length (in pixels) of the major axis, when fitting the defect to an ellipse
SAXIS	The length (in pixels) of the minor axis, when fitting the defect to an ellipse
ANGLE	The rotated angle of the LAXIS and SAXIS from vertical
SUMINT	Median of the pixel value in the defect
MEAN	Mean pixel intensity in the defect
STDEV	Standard deviation of the pixel intensity in the defect

BKMEAN	Mean of the background pixels values
BKSDEV	Standard deviation in the background pixel values
PMAX	Maximum pixel value
PMIN	Minimum pixel value
MEDIAN	Median pixel value
SIZE_MM	Size of the defect in millimeters calculated using SUMINT measurement
DIST_TO_ROI	Distance to nearest region of interest (i.e. beam footprint) from defect location
PMAX_NORM	Normalization factor used in maximum pixel value
OPTIC	Type of Optic (THG, WFL, SHG, etc)
IN_EXIT (TRUTH)	Surface of optic that defect is on (determined from microscope images of site)

```
Overall Voted Accuracy = 95.3457%
Overall Confusion Matrix:

  TRUTH
Exit Input
---- -----
 909    37   Exit
  31   484   Input      PREDICTIONS
```

Figure 5: Avatar's accuracy of prediction for a fused silica optic with over 1500 data points.

The sites that Avatar misclassified were evenly split between input and exit surfaces, leading one to think that there is no fundamental mechanism happening on one surface versus the other that Avatar could not recognize in the provided data. To make sure that our algorithm was not reporting trivial or uncorrelated patterns in the data, randomly generated patterns with a weighting of 60/40 between output and input were provided to Avatar. This study yielded a ~60% accuracy, confirming that it was in fact determining actual correlations using underlying patterns in the experimental data. It is worthwhile to note that in this study, the training data and the test data are from the same data set and in order to increase the number of data points, we have treated each observation of a site (after each laser exposure) independently so that each site contributes a multiple of data points (corresponding to the number of shots it experienced). It is possible that the machine learning is finding patterns that allow it to infer that the different data points belong to the same site, as such, correctly identifying the surface classification from the training "truth". For future work, we will train on data from one optic and apply the test on a different (but similar) optic to fully test machine learning's capabilities.

5. CONCLUSIONS

Machine learning has a demonstrated it's potential for inferring which surface a damage site is located as measured by the FODI scatter diagnostic system. It was able to successfully identify the surface of a site over 95% of the time, as opposed to analytical methods with success of less than 50%. In addition, machine learning might potentially be used in the future for comprehensive experiments to discover or "pinpoint" which parameters are important for damage prediction. Machine learning is potentially a powerful tool that can be used on other complex issues where clear analytical insight is lacking.

ACKNOWLEDGEMENT

The author would like to thank Mary Norton for her useful insight and extensive background used in this study. This work is performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

REFERENCES

1. M. A. Norton et. al., “Growth Rate of Laser Damage on the Input Surface of SiO_2 at 351 nm” in *Laser-Induced Damage in Optical Materials: 2006*, G. J. Exarhos, A. H. Guenther, K. L. Lewis, D. Ristav, M. J. Soileau, and C. J. Stolz, eds., Proc. SPIE 6403, 64030L-1-9 (2007).

2. M. A. Norton et. al., “Growth of Laser Damage in SiO_2 under Multiple Wavelength Irradiation” in *Laser-Induced Damage in Optical Materials: 2005*, G. J. Exarhos, A. H. Guenther, K. L. Lewis, D. Ristav, M. J. Soileau, and C. J. Stolz, eds., Proc. SPIE 5991, 599108-1-10 (2006).

3. A. Conder et. al., “Final Optics Damage Inspection (FODI) for the National Ignition Facility” in ” in Laser-*Induced Damage in Optical Materials: 2007*, G. J. Exarhos, A. H. Guenther, K. L. Lewis, D. Ristav, M. J. Soileau, and C. J. Stolz, eds., Proc. SPIE 6720, 672010-1-12 (2008).

4. W. Philip Kegelmeyer Jr., “The Counter-Intuitive Properties of Ensembles for Machine Learning”, Sandia Technical Report 2007-4903P (2007).

5. L. M. Kegelmeyer et. al., “NIF Optics Inspection Analysis: Overview and Update”, CASIS Workshop, Livermore, CA (2005)